

07-19-00

Express Mail Label No. EK711552729US

PTO
JCS
18 U.S.
618920

A

07/18/00

07/18/00

JCS
18 U.S.
618920
PTO**UTILITY PATENT APPLICATION TRANSMITTAL**
(Large Entity)*(Only for new nonprovisional applications under 37 CFR 1.53(b))*Docket No.
POU920000092US1

Total Pages in this Submission

TO THE ASSISTANT COMMISSIONER FOR PATENTSBox Patent Application
Washington, D.C. 20231

Transmitted herewith for filing under 35 U.S.C. 111(a) and 37 C.F.R. 1.53(b) is a new utility patent application for an invention entitled:

Job Scheduling Based Upon Availability Of Real And/Or Virtual Resources

and invented by:

Waiman Chan and Margaret M. CurryIf a **CONTINUATION APPLICATION**, check appropriate box and supply the requisite information:☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Which is a:

☐ Continuation ☐ Divisional ☐ Continuation-in-part (CIP) of prior application No.: _____

Enclosed are:

Application Elements

1. ☒ Filing fee as calculated and transmitted as described below
2. ☒ Specification having 12 pages and including the following:
 - a. ☒ Descriptive Title of the Invention
 - b. ☐ Cross References to Related Applications *(if applicable)*
 - c. ☐ Statement Regarding Federally-sponsored Research/Development *(if applicable)*
 - d. ☐ Reference to Microfiche Appendix *(if applicable)*
 - e. ☒ Background of the Invention
 - f. ☒ Brief Summary of the Invention
 - g. ☒ Brief Description of the Drawings *(if drawings filed)*
 - h. ☒ Detailed Description
 - i. ☒ Claim(s) as Classified Below
 - j. ☒ Abstract of the Disclosure

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
POU920000092US1

Total Pages in this Submission

Application Elements (Continued)

3. ☒ Drawing(s) *(when necessary as prescribed by 35 USC 113)*
- a. ☐ Formal Number of Sheets _____
- b. ☒ Informal Number of Sheets 2
4. ☒ Oath or Declaration
- a. ☒ Newly executed *(original or copy)* ☐ Unexecuted
- b. ☐ Copy from a prior application (37 CFR 1.63(d)) *(for continuation/divisional application only)*
- c. ☒ With Power of Attorney ☐ Without Power of Attorney
- d. ☐ DELETION OF INVENTOR(S)
Signed statement attached deleting inventor(s) named in the prior application,
see 37 C.F.R. 1.63(d)(2) and 1.33(b).
5. ☐ Incorporation By Reference *(usable if Box 4b is checked)*
The entire disclosure of the prior application, from which a copy of the oath or declaration is supplied
under Box 4b, is considered as being part of the disclosure of the accompanying application and is hereby
incorporated by reference therein.
6. ☐ Computer Program in Microfiche *(Appendix)*
7. ☐ Nucleotide and/or Amino Acid Sequence Submission *(if applicable, all must be included)*
- a. ☐ Paper Copy
- b. ☐ Computer Readable Copy *(identical to computer copy)*
- c. ☐ Statement Verifying Identical Paper and Computer Readable Copy

Accompanying Application Parts

8. ☒ Assignment Papers *(cover sheet & document(s))*
9. ☐ 37 CFR 3.73(B) Statement *(when there is an assignee)*
10. ☐ English Translation Document *(if applicable)*
11. ☒ Information Disclosure Statement/PTO-1449 ☒ Copies of IDS Citations
12. ☐ Preliminary Amendment
13. ☒ Acknowledgment postcard
14. ☒ Certificate of Mailing
- ☐ First Class ☒ Express Mail *(Specify Label No.):* EK711552729US

UTILITY PATENT APPLICATION TRANSMITTAL
(Large Entity)

(Only for new nonprovisional applications under 37 CFR 1.53(b))

Docket No.
POU920000092US1

Total Pages in this Submission

Accompanying Application Parts (Continued)

15. ☐ Certified Copy of Priority Document(s) *(if foreign priority is claimed)*

16. ☐ Additional Enclosures *(please identify below):*

Fee Calculation and Transmittal

CLAIMS AS FILED

For	#Filed	#Allowed	#Extra	Rate	Fee
Total Claims	4	- 20 =	0	x \$18.00	\$0.00
Indep. Claims	1	- 3 =	0	x \$78.00	\$0.00
Multiple Dependent Claims (check if applicable) <input type="checkbox"/>					\$0.00
BASIC FEE					\$690.00
OTHER FEE <i>(specify purpose)</i>					\$0.00
TOTAL FILING FEE					\$690.00

- ☐ A check in the amount of _____ to cover the filing fee is enclosed.
- ☒ The Commissioner is hereby authorized to charge and credit Deposit Account No. **09-0463** as described below. A duplicate copy of this sheet is enclosed.
- ☒ Charge the amount of **\$690.00** as filing fee.
 - ☒ Credit any overpayment.
 - ☒ Charge any additional filing fees required under 37 C.F.R. 1.16 and 1.17.
 - ☐ Charge the issue fee set in 37 C.F.R. 1.18 at the mailing of the Notice of Allowance, pursuant to 37 C.F.R. 1.311(b).


Signature

Lawrence D. Cutter - Attorney
Registration No.: 28,501
IBM Corporation
Intellectual Property Law
2455 South Road, M/S P386
Poughkeepsie, New York 12601
Telephone: (845) 433-1172
Facsimile: (845) 432-9601

Dated: **July 18, 2000**

cc:

Docket No : POU920000092US1

Inventor : W. Chan et al.

Title : JOB SCHEDULING BASED UPON
AVAILABILITY OF REAL
AND/OR VIRTUAL RESOURCES

JC813 U.S. PTO
09/618920
07/18/00

APPLICATION FOR UNITED STATES
LETTERS PATENT

"Express Mail" Mailing Label No.: EK711552729US
Date of Deposit: July 18, 2000

I hereby certify that this paper is being
deposited with the United States Postal Service
as "Express Mail Post Office to Addressee" service
under 37 CFR 1.10 on the date indicated above
and is addressed to: Box Patent Application,
Assistant Commissioner for Patents, Washington,
D.C. 20231.

Name: BILLY R. STACY

Signature: Billy R. Stacy

INTERNATIONAL BUSINESS MACHINES CORPORATION

JOB SCHEDULING BASED UPON AVAILABILITY OF REAL AND/OR VIRTUAL RESOURCES

Background of the Invention

5 The present invention is generally directed to systems and methods for scheduling jobs running on a data processing system. More particularly, the present invention is directed to providing user control of a wide range of computer related resources which exists at a plurality of levels within a data processing system. Even more particularly, the present invention is directed to the utilization of a consistent set of resource models extending over different data processing levels from a low hardware level through the operating system level and onto an administrative and user level.

10 With the proliferation of data processing systems employing multiple processors it has typically been the role of the operating system to schedule jobs based upon job descriptions provided by system users. In contrast, the present invention breaks away from this modality by providing mechanisms for much greater user involvement in the utilization of data processing system resources.

15 In particular, it is noted that the increased use of data processing systems employing multiple processors particularly in a symmetric multiprocessing (SMP) environment, there has been provided an opportunity for various tasks associated with a user's job to be run on selected ones of a number of processing units. Accordingly, the present invention includes a mechanism for user controlled association of various job tasks with particular processors. Accordingly, this permits greater control of the data processing operation by an entity which is most familiar with the structure of the program or application, namely the user of application software. For example, users running application programs which are particularly susceptible to parallelization are provided with a specific mechanism for linking various job tasks with different processors.

While it is noted that the present invention is particularly directed to symmetric multi processing environments and to the allocation of processors as a data processing system resource, the present invention also extends to a more generic modality in which a number of different system resources are modeled and defined in a consistent way across a number of different system levels.

In particular, it is noted that resources such as the amount of memory available may be defined as being available to a user as a hardware level resource. It is also noted that resources such as virtual memory or the number of CPUs available may be defined and modeled at the operating system level so as to provide a mechanism for user control of these resources in a direct fashion. Additionally, there are resources available at the administrative level, a level which overlies the operating system and is available in such systems as the assignee's IBM's RS/6000 SP product employing LoadLeveler as part of this administrative level. Resources now available at this administrative level include such items as available software licensing. This might occur for example in a situation in which a desired licensed application program has associated with it a contractual obligation that the program may not be run on more than N processors at the same time under the given license. Likewise it may be the case that a software license is limited by the requirement that the software run only on specified (and/or identified) processing units.

Additionally, users wanting to run large simulation programs often need to be able to keep all of the relevant simulation data in memory at the same time in order for the application to run efficiently. However, the allocation of memory as a resource has typically only been accomplished through the operating system, and has not been under user control. Furthermore, some jobs actually have a requirement that more than one CPU is involved in the computation.

In addition to placing all of these computer resources within greater control of a user, it is also very desirable that any such scheduling is done in a dynamic configuration without any requirements for rebooting or otherwise intervening with operating system functions.

POU920000092US1

In order to best understand the proper context in which the present invention operates, it is noted that the present invention does not operate at the microsecond or nanosecond time scales of the operating system. Rather, the present invention operates as a part of a macro scheduler operating above the operating system level. Nonetheless through the utilization of API-like (Application Program Interface) statements, proper communication with the operating system level is established in a consistent manner using a resource model which is applicable over several data processing system levels. In particular, the model employed by the present invention are such that resources are specified in terms of both type and amounts. Thus, there is typically a quantity or capacity associated with each one of the resources which are modeled in the present invention.

Summary of the Invention

In accordance with a preferred embodiment of the present invention a method for managing resources in a data processing system comprises an initial step of defining at least one type of resource together with a quantity which is associated with that resource type. This quantity is indicative of resource capacity. The resource is either a hardware level resource, an operating system level resource, an application level resource an a user level resource. A determination is made as to whether a user has requested the use of one of these resources. Then the availability of the requested resource is determined with respect to its indicated type and quantity. Finally a user job is dispatched in accordance with the available resource type and quantity.

The present invention functions by providing a consistent resource model which extends across multiple data processing systems levels (hardware through operating system through administrative and finally through user levels). The resource model provided herein is defined in terms of type and quantity. The model is invoked through the utilization of API like statements provided by the user (or their interactive equivalents). These statements are analyzed in order to determine the requested resource usage. These analyzed statements are then used to schedule and to run jobs which use the specified resources. The power and flexibility of the present invention

derives from the fact that a range of resource types is modeled in a consistent fashion across a number of data processing system levels particularly with respect to the incorporation of type and quantity specifications and in the modeling of a large variety of resources which are contemplated as being handled in accordance with the present method.

5 Accordingly, it is an object of the present invention to provide users with greater control of how their applications are run.

It is also an object of the present invention to provide a better match between application programs and the resources which they require in a data processing system complex.

10 It is also an object of the present invention to permit a user to control the number of processors employed in the running of the user's application program particularly in those environments in which there are a plurality of processors such as in a symmetric multi processor (SMP) architecture.

It is also an object of the present invention to permit users to control the type and amount of resources that a job consumes whether or not the resource is real or virtual.

15 It is a still further object of the present invention to provide a mechanism for accommodating constraints associated with limited software program licenses.

It is a still further object of the present invention to permit a user to control the amount of real or virtual memory used in processing a job or task.

20 It is also an object of the present invention to control the utilization of hardware resources by a user in a dynamic fashion which does not require either rebooting or operating system intervention.

It is a still further object of the present invention to provide a consistent resource model operable over a plurality of layers in a digital processing computer system or complex.

It is also an object of the present invention to improve the utilization of symmetric multi processing system environments.

5 It is also an object of the present invention to permit real time modification of resource requirements with respect to both type and to quantity.

Lastly, but not limited hereto, it is an object of the present invention to provide a mechanism for recovery of allocated and requested resources upon job and/or task completion.

10 The recitation herein of a list of desirable objects which are met by various embodiments of the present invention is not meant to imply or suggest that any or all of these objects are present as essential features, either individually or collectively, in the most general embodiment of the present invention or in any of its more specific embodiments.

Description of the Drawings

15 The subject matter which is regarded as the invention is particularly pointed out and distinctly claimed in the concluding portion of the specification. The invention, however, both as to organization and method of practice, together with the further objects and advantages thereof, may best be understood by reference to the following description taken in connection with the accompanying drawings in which:

20 Figure 1 is a block diagram illustrating the embodiment of the environment in which the present invention is employed and particularly illustrating its utilization across several hierarchical layers found in data processing systems; and

Figure 2 is a block diagram illustrating the resource model employed in the present invention and more particularly illustrating the utilization of that model in its relationship to various parts of the system.

Detailed Description of the Invention

5 Figure 1 illustrates the environment in which the present invention is employed. In particular, Figure 1 illustrates a data processing environment which includes a plurality of levels. At the lowest point there is hardware level 100. On top of level 100 there is a operating system level 200. Additionally, there is a administrative level 300 which interacts with customer, user or job level 400. Applicants submit software applications to the data processing system through
10 means of administrative level 300. In particular, in preferred embodiments of the present invention administrative level 300 is provided by a product offered by applicants assignee called LoadLeveler. LoadLeveler interacts with an applicant through the utilization of various Application Program Interfaces (APIs). These APIs may be invoked by an application or user either via keyboard commands or alternatively via interactive movements of another input/output
15 device such as a mouse, joy stick, stylus, light pen, graphics tablet or the like. Administrative level 300 includes scheduler 301 which employs a number of class definitions 302 with which scheduler 301 interacts, typically by means of internal APIs which are not available to the user in a direct fashion. Operating system level 200 typically manages resources such as memory 201 particularly in the form of virtual memory and also manages the CPU resource availability.
20 Likewise, at hardware level 100, resources include such things as memory 101 disk storage or DASD 102, tape drive 104 and CPUs 103. The CPU management at operating system level 200 and hardware level 100 are not the same. This is why a resource such as CPUs is present at both the hardware and operating system levels.

25 Consideration is now given to a discussion of the resource model used in the present invention which is described in Figure 2. Figure 2 illustrates the model in terms of the relation between various objects in the resource model employed herein. As a general comment the lines drawn between the various software component objects shown illustrate a "has" relationship.

Furthermore the numbers adjacent to the boxes and the lines indicate the nature of the relationship. In particular it may be a one-to-one relation or in other cases it may be a one-to-N relationship where N sometimes varies between 1 and n and other times it varies between 0 and n. In particular, job step 500 has one or more tasks. Each job step 500 has one step identifier 510 (or stepid meaning "step identifier "or" step ID"). Each task 501 has 0 or more resource requirements 502. Each resource requirement in accordance with the present model is characterized by both a type and amount 503. A machine 520 has zero or more resources 525. And again each resource rather is characterized with type and amount information 526. Cluster 530 has one or more machines 520. Each resource is characterized by both a type and an amount. Resource 525 has 0 or more resource usage's 511. Each resource usage is characterized by amount 512. And a resource usage has one step identifier (stepid) 510.

Resources are defined at the discretion of the administrator, representing either physical characteristics of a machine or virtual characteristics meaningful for the system and the workload represented. The resources are defined in accordance with available in a specified quantity on a specific machine or are defined to be available aggregately or on any of the machines defined to the system. For example, resources may be defined to represent any of the following identities:

1. The number of processors in a symmetric multi processing (SMP) machine;
2. The amount of real memory available on a machine;
3. The amount of virtual memory defined on a machine;
4. The number of installed software licenses available on a machine;
5. The number of installed floating software licenses available aggregately or on the set of machines defined to the system; and

6. Any other named, quantifiable entity that is meaningful to the scheduling of jobs on machines that are defined to the system.

Scheduling of jobs is based on their requested resource consumption with scheduling directives being provided in a flexible manner. In particular, users specify the type and amount of resource that each job consumes when running. To facilitate parallel applications resource specification is made based on the amount of each indicated resource which is anticipated to be consumed by each process of the user's application program. Administrators specify which resources (of those defined) are to be considered when scheduling jobs. A job is scheduled to run on machine(s) based on two factors:

1. scheduling directives indicating which defined resources to consider;

2. sufficiency of those resources based upon the job's requests to consume a specified amount of these resources per each of its processes or tasks. Availability of resources is constrained by consumption of resources by currently running jobs.

Consumption of resources during job execution and the restoration of resource availability upon job completion is also provided. In particular, resources are said to be consumed by a job and are considered to be exclusively used by the job during its execution. This renders these resources unavailable for consumption by another job unless or until the job consuming the resource completes or is otherwise canceled by the user or terminated by the system. At that time, the resources consumed by that job are restored to the system and become available for future consumption by other jobs which are scheduled to run.

The accessibility to real time resource availability information is also provided in the present invention. In particular, resource availability information is made accessible via APIs so that an external scheduler or other resource monitoring an application may utilize that information. In addition, users may access resource availability via commands. Availability information is provided in a real time fashion so as to represent resources available at the time

that the request for information is made regardless of the means of that request, taking into consideration the resources initially defined and those currently consumed by actively running jobs.

Adaptability for changing resource definitions is also made available as a means for affecting job scheduling. In particular, the method of the present invention accommodates real time modification of the resources and the quantities associated therewith such that the scheduler adapts to changes in availability while still continuing to attempt to schedule jobs. This reconfiguration process allows the increase or decrease of available resources regardless of the amounts of those resources which are currently defined to be available and which are considered to be consumed by running jobs. Thus resources may be removed from the configuration or their available capacity levels are reduced below that which is currently in use by running jobs.

Recovery capability is also provided by the present invention. For example in the event that there is a failure of the job scheduling system such that a predefined alternate scheduler takes over for one that has failed, currently executing jobs send information to the new scheduler indicating their resource consumption. This prevents the new scheduler from scheduling jobs that consume resources which are actually in use by already actively running jobs.

From the above comments it should be appreciated that the model of the resource objects as provided by the present invention as illustrated herein provides all of the advantages listed above. It is further seen that the objects of the present invention, as described above, are also met. In particular, it is seen that the present invention puts much greater flexibility and control in the hands of an application program, author and/or user.

While the invention has been described in detail herein in accordance with certain preferred embodiments thereof, many modifications and changes therein may be effected by those skilled in the art. Accordingly, it is intended by the appended claims to cover all such modifications and changes as fall within the true spirit and scope of the invention.

Claims:

What is claimed is:

1. A method for managing resources in a data processing system, said method comprising:
 - defining at least one type of resource, together with a quantity associated with that type of resource, said quantity being indicative of resource capacity, said type of resource being selected from the group consisting of hardware level resources, operating system level resources and application level resources;
 - determining whether a user has requested use of said at least one resource;
 - determining availability of said requested resource as to type and quantity;
 - dispatching a user job which requests said resource upon the condition that said resource is available in the type and quantity requested.
2. The method of claim 1 in which said hardware level resources are selected from the group consisting of CPUs and random access memory.
3. The method of claim 1 in which said operating system level resource is virtual memory.
4. A method for providing a user with control of a data processing system, said method comprising the steps of:
 - analyzing user supplied command statements which provide access to resources having internal resource models, said resource models including, for each resource, a type and a quantity description, wherein at least one resource model is defined consistently across multiple data processing system levels; and

[illegible]

5 A model for computer resources based on type and amount is employed in a consistent manner across several layers in a data processing system as a mechanism for providing users with significantly more control over the allocation, utilization and scheduling of resources for user's software applications running on the target computer system.

POU920000092US1

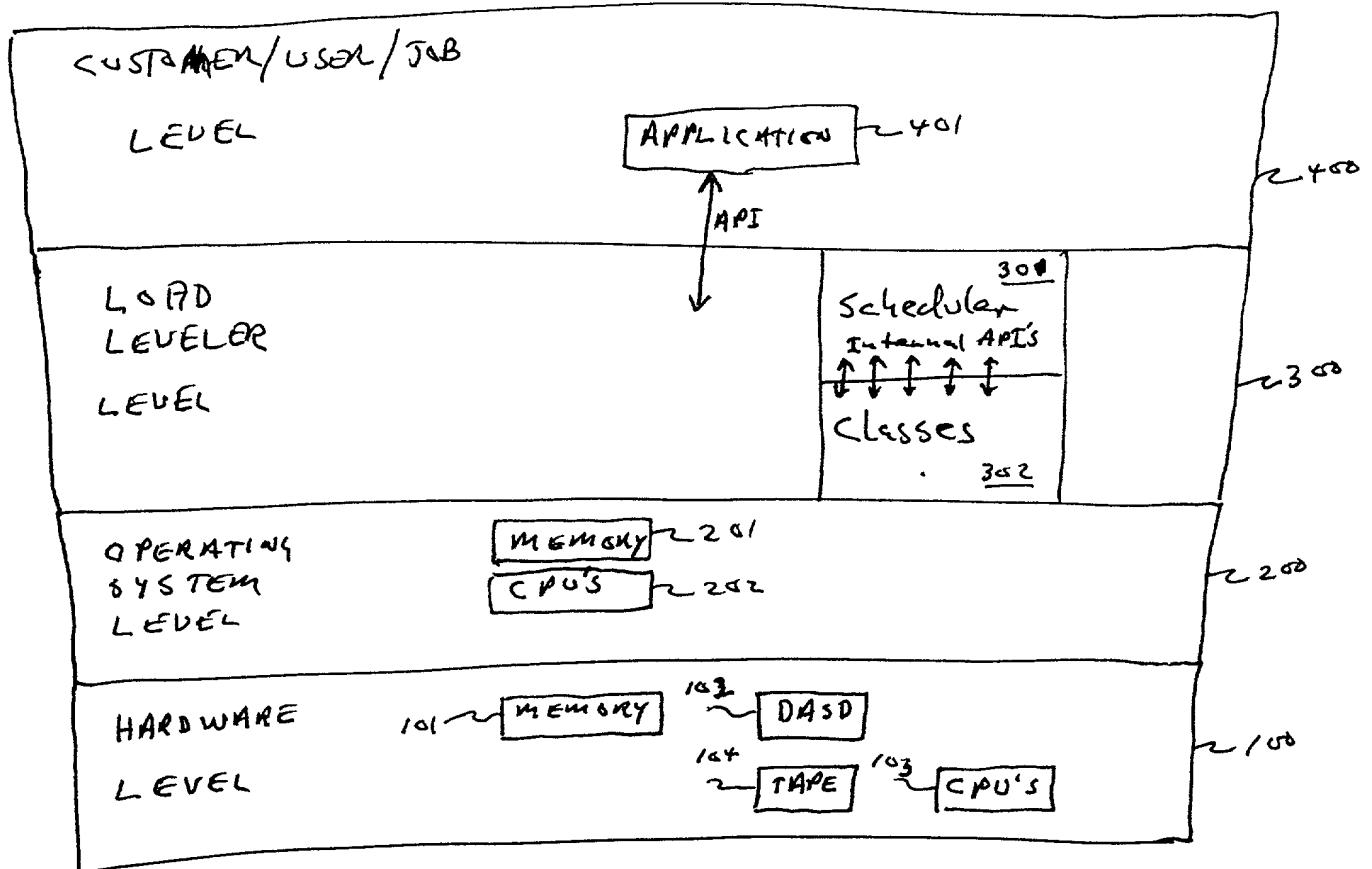


Figure 1




Figure 2

Docket No.
POU920000092US1

Declaration and Power of Attorney For Patent Application

English Language Declaration

As a below named inventor, I hereby declare that:

My residence, post office address and citizenship are as stated below next to my name,

I believe I am the original, first and sole inventor (if only one name is listed below) or an original, first and joint inventor (if plural names are listed below) of the subject matter which is claimed and for which a patent is sought on the invention entitled

JOB SCHEDULING BASED UPON AVAILABILITY OF REAL AND/OR VIRTUAL RESOURCES

the specification of which

(check one)

☒ is attached hereto.

☐ was filed on _____ as United States Application No. or PCT International Application Number _____ and was amended on _____ (if applicable)

I hereby state that I have reviewed and understand the contents of the above identified specification, including the claims, as amended by any amendment referred to above.

I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, Code of Federal Regulations, Section 1.56.

I hereby claim foreign priority benefits under Title 35, United States Code, Section 119(a)-(d) or Section 365(b) of any foreign application(s) for patent or inventor's certificate, or Section 365(a) of any PCT International application which designated at least one country other than the United States, listed below and have also identified below, by checking the box, any foreign application for patent or inventor's certificate or PCT International application having a filing date before that of the application on which priority is claimed.

Prior Foreign Application(s)

Priority Not Claimed

_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>
_____ (Number)	_____ (Country)	_____ (Day/Month/Year Filed)	<input type="checkbox"/>

I hereby claim the benefit under 35 U.S.C. Section 119(e) of any United States provisional

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

(Application Serial No.)

(Filing Date)

I hereby claim the benefit under 35 U. S. C. Section 120 of any United States application(s), or Section 365(c) of any PCT International application designating the United States, listed below and, insofar as the subject matter of each of the claims of this application is not disclosed in the prior United States or PCT International application in the manner provided by the first paragraph of 35 U.S.C. Section 112, I acknowledge the duty to disclose to the United States Patent and Trademark Office all information known to me to be material to patentability as defined in Title 37, C. F. R., Section 1.56 which became available between the filing date of the prior application and the national or PCT International filing date of this application:

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

(Application Serial No.)

(Filing Date)

(Status)
(patented, pending, abandoned)

I hereby declare that all statements made herein of my own knowledge are true and that all statements made on information and belief are believed to be true; and further that these statements were made with the knowledge that willful false statements and the like so made are punishable by fine or imprisonment, or both, under Section 1001 of Title 18 of the United States Code and that such willful false statements may jeopardize the validity of the application or any patent issued thereon.

POWER OF ATTORNEY: As a named inventor, I hereby appoint the following attorney(s) and/or agent(s) to prosecute this application and transact all business in the Patent and Trademark Office connected therewith. *(list name and registration number)*

William B. Porter, Reg. No. 33,135

Floyd A. Gonzalez, Reg. No. 26,732

Lynn L. Augspurger, Reg. No. 24,227

William A. Kinnaman, Jr., Reg. No. 27,650

Lily Neff, Reg. No. 38,254

Marc A. Ehrlich, Reg. No. 39,966

Lawrence D. Cutter, Reg. No. 28,501

Christopher A. Hughes, Reg. No. 26,914

Edward A. Pennington, Reg. No. 32,588


John E. Hoel, Reg. No. 26,279


Joseph C. Redmond, Reg. No. 18,753

Andrew J. Wojnicki, Jr., Reg. No. 43,995

Send Correspondence to: **Lawrence D. Cutter**
IBM Corporation, IPLAW
2455 South Road, M/S P386
Poughkeepsie, NY 12601

Direct Telephone Calls to: *(name and telephone number)*
Lawrence D. Cutter Telephone: (845) 433-1172

Full name of sole or first inventor Waiman Chan
Sole or first inventor's signature  Date 7/18/2000
Residence 50 Hagan Drive, Poughkeepsie, NY 12603
Citizenship United States of America
Post Office Address 50 Hagan Drive, Poughkeepsie, NY 12603

Full name of second inventor, if any Margaret M. Curry
Second inventor's signature  Date 7/18/2000
Residence 119 Pinebrook Drive, Hyde Park, NY 12538
Citizenship United States of America
Post Office Address 119 Pinebrook Drive, Hyde Park, NY 12538